

# Outcome prediction in mathematical models of immune response to infection

Manuel Mai<sup>1,\*</sup>, Kun Wang<sup>2,3</sup>, Greg Huber<sup>4</sup>, Michael Kirby<sup>2,5</sup>, Mark D. Shattuck<sup>6,3</sup>,  
Corey S. O'Hern<sup>1,3,7,8</sup>

**1** Department of Physics, Yale University, New Haven, CT, USA

**2** Department of Mathematics, Colorado State University, Fort Collins, CO, USA

**3** Department of Mechanical Engineering and Material Science, Yale University, New Haven, CT, USA

**4** Kavli Institute for Theoretical Physics, Kohn Hall, University of California, Santa Barbara, CA, USA

**5** Department of Computer Science, Colorado State University, Fort Collins, CO, USA

**6** Benjamin Levich Institute and Physics Department, The City College of New York, New York, NY, USA

**7** Department of Applied Physics, Yale University, New Haven, CT, USA

**8** Graduate Program in Computational Biology and Bioinformatics, Yale University, New Haven, CT, USA

\* E-mail: manuel.mai@yale.edu

## Abstract

Clinicians need to predict patient outcomes with high accuracy as early as possible after disease inception. In this manuscript, we show that patient-to-patient variability sets a fundamental limit on outcome prediction accuracy for a general class of mathematical models for the immune response to infection. However, accuracy can be increased at the expense of delayed prognosis. We investigate several systems of ordinary differential equations (ODEs) that model the host immune response to a pathogen load. Advantages of systems of ODEs for investigating the immune response to infection include the ability to collect data on large numbers of ‘virtual patients’, each with a given set of model parameters, and obtain many time points during the course of the infection. We implement patient-to-patient variability  $v$  in the ODE models by randomly selecting the model parameters from Gaussian distributions with variance  $v$  that are centered on physiological values. We use logistic regression with one-versus-all classification to predict the discrete steady-state outcomes of the system. We find that the prediction algorithm achieves near 100% accuracy for  $v = 0$ , and the accuracy decreases with increasing  $v$  for all ODE models studied. The fact that multiple steady-state outcomes can be obtained for a given initial condition, *i.e.* the basins of attraction overlap in the space of initial conditions, limits the prediction accuracy for  $v > 0$ . Increasing the elapsed time of the variables used to train and test the classifier, increases the prediction accuracy, while adding explicit external noise to the ODE models decreases the prediction accuracy. Our results quantify the competition between early prognosis and high prediction accuracy that is frequently encountered by clinicians.

## Introduction

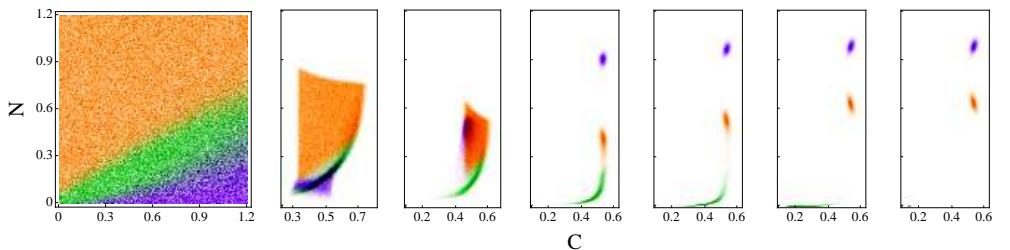
The immune response to infection is a complex process that involves a wide range of length scales from proteins to cells [1–4], tissues [5], and organ systems [6]. Despite enormous progress over the past 30 years in developing mathematical models for the immune response to infectious disease such as tuberculosis [7], HIV [8–11], and influenza [12,13], these models still have not been able to dramatically improve patient diagnosis, prognosis, and treatment [14,15]. Instead, vaccine and drug development often relies on costly trial-and-error methods [16]. However, advances in gene sequencing capabilities [17], increasing speeds of computer processors, and the ability to store enormous amounts of medical data promise dramatic improvements in mathematical approaches to predicting and controlling the response to infectious disease [18–20].

One promising mathematical approach is to use machine learning methods on large data sets to classify patients as healthy or sick, perform early warning analyses for early detection of infection, or identify the minimal set of genes responsible for a particular immune response. [21,22] However, many questions are left unanswered in such studies. For example, how much and what kinds of data are required to have confidence in the machine learning predictions and what are the underlying biophysical mechanisms for the relationships between variables that are identified by these techniques? Further, it is difficult to determine differences in the immune response that arise from patient-to-patient variations compared to slight differences in the initial conditions of each patient.

In this manuscript, we focus on sets of ordinary differential equations (ODEs) as mathematical models for the immune response to infection. The advantages of ODEs are manifold: 1) Each ‘virtual patient’ can be considered as a set of parameters in the set of ODEs; 2) There is essentially no limit on the amount of data that can be collected on each virtual patient; 3) The accuracy of machine learning predictions can be explicitly tested as a function of the number of time points and initial conditions for each patient and the number of patients included in the training and testing sets; and 4) analysis of the fixed points (or steady-state outcomes) and basins of attraction of the ODEs can give biophysical insight into the immune response to infection.

We will investigate several classes of ODE models for the immune response to infection. First, we will describe a four-dimensional model for the acute inflammatory response to a pathogen load that was studied in detail in Ref. [23]. We will then consider reduced versions of this model with fewer variables and parameters obtained by slaving one or more of the original four variables, as well as changes to form of the ODEs that alter the fixed point structure and flows between them. For each model, a virtual patient is defined by one set of parameters. Given an initial condition (values of the variables at time  $t = 0$ ), the patient will evolve deterministically to one of several possible discrete steady-state ( $t \rightarrow \infty$ ) outcomes, or fixed points. Thus, for each patient, we can determine the basins of attraction that map initial conditions for all of the variables to steady-state outcomes by numerically integrating the sets of ODEs.

We seek to determine the limits of the prediction accuracy of discrete steady-state outcomes of ODEs as a function of patient variability (*i.e.* random fluctuations in parameter values) using machine learning techniques. In the limit of zero patient variability, our simple classification algorithm (logistic regression) can achieve nearly perfect prediction accuracy even when the classification occurs on variables at short times. However, as the patient variability increases, the basins of attraction for different patients yield different outcomes for a given initial condition as shown in Fig. 1 for 2% patient variability in model (1) for the immune response to infection. (See Materials and Methods.) The fact that each set of initial conditions does not possess a unique outcome places a fundamental limit on the predictability of patient



**Figure 1. Time evolution of patient outcomes for a range of neutrophil ( $N$ ) and cortisol ( $C$ ) initial conditions for Model A.** (left) Patient outcomes in the long-time limit given  $N$  and  $C$  initial conditions for model (1) are shaded green, orange, and purple for the steady-state outcomes of health, aseptic, and septic death, respectively. The initial values of the pathogen load and damage are  $P_0 = 0.35$  and  $D_0 = 0$ , and patient variability is set to  $v = 2\%$ . The right six panels indicate how the systems in the leftmost panel separate in the  $N$  and  $C$  plane as time increases,  $t = 10, 20, 50, 100, 250$ , and  $500$ , from left to right.

outcomes. Thus, we find that the machine learning prediction accuracy decreases with increasing patient variability. In contrast, for a given patient variability, the prediction accuracy increases with the time used for classification as the systems converge to their steady-state outcomes (Fig. 1). We also show that at short times our classification algorithm saturates the theoretical limit for the prediction accuracy in the presence of patient variation, and that the addition of external noise only worsens the outcome prediction accuracy.

The manuscript is organized as follows. In the Materials and Methods section, we introduce several ODE systems that have been used to model the host immune response to infection [23], including their parameter sets and discrete steady-state outcomes, and describe how we implement patient variability in the ODE models. In the Results section, we emphasize our three main results that hold for all of the ODE models we studied: 1) patient variability leads to overlap of the basins of attraction for the steady-state outcomes, which limits the outcome prediction accuracy, 2) the prediction accuracy increases with the time used for classification because the basins of attraction separate with increasing time, and 3) the addition of external measurement noise further reduces the prediction accuracy. In the Discussion section, we point out the clinical implications of our work and describe important future studies of the prediction accuracy for ODE models with continuous outcomes. In the Supporting Information, we discuss a generalized one-dimensional ODE against which we compare our results and the numerical implementation of the algorithms used to classify the steady-state outcomes of the ODEs.

## Materials and Methods

Our studies focus on several ODE models with varying complexity for the immune response to pathogen load that were first introduced in Ref. [23]. These ODE models can have up to four coupled variables that represent the concentration of pathogen  $P$ , activated neutrophils  $N$ , inflammation (or damage)  $D$ , and immuno-suppressor (cortisol)  $C$ . The models include interactions between these four quantities. For example, the presence of pathogen  $P > 0$  causes an immune response, where neutrophils are activated and  $N$  increases. Neutrophils kill pathogen, which decreases  $P$ , but also cause inflammation (damage), which increases  $D$ . The cortisol level  $C$  increases when there is a high neutrophil level, which then reduces the neutrophil level.

Model (1) (Eqs. 1-4) includes all four variables  $P$ ,  $N$ ,  $D$ , and  $C$ . The right-hand side of  $dP/dt$  is a sum of three terms. The first term enables logistic growth of the pathogen. In the absence of any other terms, any positive initial  $P_0$  will cause  $P$  to grow logically to the steady-state value  $P_\infty$ . The second term mimics a local, non-specific response to an infection. For small values of  $P$ , the decrease is proportional to  $P$ . For larger values of  $P$ , the decrease caused by the second term is constant. The third term models the decrease of  $P$  due to interactions with activated immune cells (neutrophils)  $N$ . Activated neutrophils  $N$  can directly decrease  $P$ . The anti-inflammatory response, which is captured by the cortisol level  $C$ , mitigates this effect leading to a decrease of  $P$  proportional to  $N * P / (1 + (C/C_\infty)^2)$ .

Two terms determine the rate of change in neutrophils,  $dN/dt$ . The first term accounts for the fact that neutrophils can be activated if a resting neutrophil cell encounters a pathogen  $P$  or an already activated neutrophil  $N$ . Furthermore, tissue damage  $D$  also triggers the activation of neutrophils. The second term describes the death of neutrophils  $N$ , with the decrease in  $N$  proportional to the amount of neutrophils present.

The rate of change in damage  $dD/dt$  is also controlled by two terms. The first term mimics positive feedback between  $D$  and  $N$ . Activated phagocytes cause collateral damage in the tissue. Again, the effectiveness of  $N$  is mitigated by the anti-inflammatory response  $1/(1 + (C/C_\infty)^2)$ . The saturation function  $f_s$  models the fact that the effect of  $N$  on  $D$  saturates for large  $N$ . The second term,  $-\mu_d D$ , represents repair of the tissue.

The anti-inflammatory response  $C$  increases with the source term  $s_c$ . In addition, there is a natural death rate  $\mu_c$ , which leads to a positive steady-state value of  $C$  in the absence of any immune activation  $N$  or damage  $D$ . However, even small amounts of damage and neutrophils will up-regulate  $C$ . In the case of small  $N + k_{cnd}D$ , the production of  $C$  is proportional to  $N + k_{cnd}D$ , while for large values of  $N + k_{cnd}D$ , changes in  $C$  are proportional to  $k_{cn}$ . Again, the effectiveness of  $N$  is mitigated by  $1/(1 + (C/C_\infty)^2)$ .

Model (1) has 21 parameters:  $k_{pm}$ ,  $k_{mp}$ ,  $s_m$ ,  $\mu_m$ ,  $k_{pg}$ ,  $P_\infty$ ,  $k_{pn}$ ,  $k_{np}$ ,  $k_{nn}$ ,  $s_{nr}$ ,  $\mu_{nr}$ ,  $\mu_n$ ,  $k_{nd}$ ,  $k_{dn}$ ,  $x_{dn}$ ,  $\mu_d$ ,  $C_\infty$ ,  $s_c$ ,  $k_{cn}$ ,  $k_{cnd}$ , and  $\mu_c$ . Depending on the values of these parameters, model (1) possesses different numbers of fixed points with varying stabilities. However, we will focus on a specific parameter regime (given in Table 1) with three stable fixed points, which correspond to the physiological steady-state outcomes: health, septic death, and aseptic death.

### Model (1)

$$\frac{dP}{dt} = k_{pg}P \left(1 - \frac{P}{P_\infty}\right) - \frac{k_{pm}s_mP}{\mu_m + k_{mp}P} - k_{pn}f(N)P \quad (1)$$

$$\frac{dN}{dt} = \frac{s_{nr}R}{\mu_{nr} + R} - \mu_nN \quad (2)$$

$$\frac{dD}{dt} = k_{dn}f_s(f(N)) - \mu_dD \quad (3)$$

$$\frac{dC}{dt} = s_c + \frac{k_{cn}f(N + k_{cnd}D)}{1 + f(N + k_{cnd}D)} - \mu_cC, \quad (4)$$

where

$$\begin{aligned} R &= f(k_{nn}N + k_{np}P + k_{nd}D), \\ f(V) &= V/(1 + (C/C_\infty)^2), \\ f_s(V) &= V^6/(x_{nd}^6 + V^6). \end{aligned}$$

**Table 1. Parameter mean values for model (1)**

Parameter $q$	Mean Value $\mu_q$
$k_{pm}$	0.6
$k_{mp}$	0.01
$s_m$	0.005
$\mu_m$	0.002
$k_{pg}$	0.6
$P_\infty$	20.0
$k_{pn}$	1.8
$k_{np}$	0.1
$k_{nn}$	0.01
$s_{nr}$	0.08
$\mu_{nr}$	0.12
$\mu_n$	0.05
$k_{nd}$	0.02
$k_{dn}$	0.35
$x_{dn}$	0.06
$\mu_d$	0.02
$C_\infty$	0.28
$s_c$	0.0125
$k_{cn}$	0.04
$k_{cnd}$	48.0
$\mu_c$	0.1

Models (2)-(5) given below are simplified versions of model (1). A summary of the dimension, number of parameters, and number of stable fixed points for each of the ODE models is shown in Table 2. To obtain model (2) from (1),  $C$  is set to a constant  $\bar{C} = 0.23$  and the remaining terms define a three-variable model with  $P$ ,  $N$ , and  $D$ . For model (3), we set  $C = \bar{C}$  and  $D = 0$ , which gives a two-variable model for  $P$  and  $N$ . For model (4), we set  $C = \bar{C}$  and  $P = 0$  to obtain a two-variable model for  $N$  and  $D$ . In this model, the value of the initial rise in  $N$  can be thought of as the response to trauma. For model (5), we set  $C = \bar{C} = 0.1$ ,  $D = 0$ , and  $N = 0$ , which gives a one-dimensional model for  $P$ . This model only treats the innate immune response with no activated neutrophils.

**Table 2. Summary of ODE models**

Model	Dimension	Parameters q	Stable Fixed Points	$k_{pg}$	$\bar{C}$
1	4	21	3	0.6	N/A
2	3	18	3	1.2	0.23
3	2	14	2	1.2	0.23
4	2	8	2	1.2	0.23
5	1	6	2	0.6	0.1
6	1	2	2	N/A	N/A

### Model (2)

$$\frac{dP}{dt} = k_{pg}P \left(1 - \frac{P}{P_\infty}\right) - \frac{k_{pm}s_mP}{\mu_m + k_{mp}P} - k_{pn}f(N)P \quad (5)$$

$$\frac{dN}{dt} = \frac{s_{nr}R}{\mu_{nr} + R} - \mu_nN \quad (6)$$

$$\frac{dD}{dt} = k_{dn}f_s(f(N)) - \mu_dD \quad (7)$$

$$(8)$$

### Model (3)

$$\frac{dP}{dt} = k_{pg}P \left(1 - \frac{P}{P_\infty}\right) - \frac{k_{pm}s_mP}{\mu_m + k_{mp}P} - k_{pn}f(N)P \quad (9)$$

$$\frac{dN}{dt} = \frac{s_{nr}R_3}{\mu_{nr} + R_3} - \mu_nN, \quad (10)$$

$$(11)$$

where

$$R_3 = f(k_{nn}N + k_{np}P).$$

### Model (4)

$$\frac{dN}{dt} = \frac{s_{nr}R_4}{\mu_{nr} + R_4} - \mu_nN \quad (12)$$

$$\frac{dD}{dt} = k_{dn}f_s(f(N)) - \mu_dD, \quad (13)$$

$$(14)$$

where

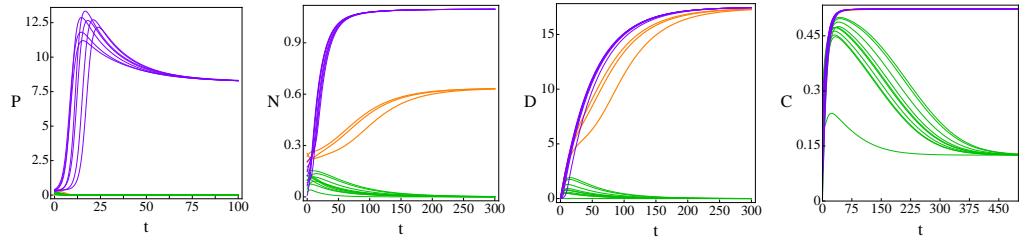
$$R_4 = f(k_{nn}N + k_{nd}D).$$

### Model (5)

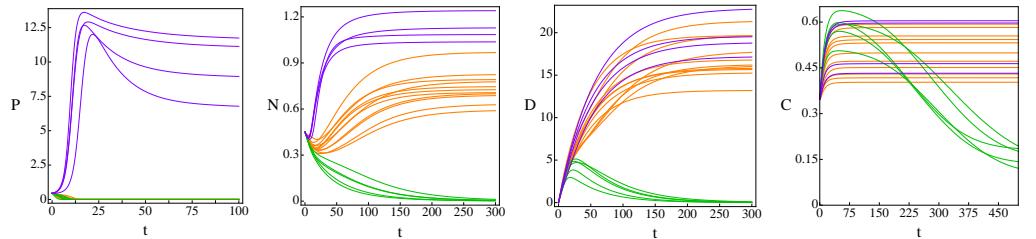
$$\frac{dP}{dt} = k_{pg}P \left(1 - \frac{P}{P_\infty}\right) - \frac{k_{pm}s_mP}{\mu_m + k_{mp}P} \quad (15)$$

In Fig. 2, we show the time evolution of the four variables  $P$ ,  $N$ ,  $D$ , and  $C$  for model (1) for twenty different sets of random initial conditions to illustrate its three stable fixed points (health, septic death, and aseptic death) using the parameter values in Table 1. For trajectories that approach the septic death fixed point, the pathogen and neutrophil levels grow rapidly. The high neutrophil level causes cortisol to increase as well. Despite the high level, the neutrophils cannot reduce the pathogen load and the cortisol level is not large enough to reduce the neutrophil level. As a result, the high neutrophil level causes significant damage at long times, which is termed septic death due to the associated high pathogen level. Thus, the septic death steady-state outcome is characterized by  $P > 0$ ,  $N > 0$ ,  $D > 0$ , and  $C > C_\infty$ .

In the healthy state, the pathogen level can be reduced to zero by the neutrophils, and the neutrophil level can be reduced to zero by cortisol. Once the neutrophil level is zero, the cortisol level returns to its background level and damage decreases to zero. Thus, the healthy state is characterized by  $P = 0$ ,  $N = 0$ ,  $D = 0$ , and  $C = C_\infty$ .



**Figure 2.**  $P$ ,  $N$ ,  $D$ , and  $C$  versus time for model (1) from 20 random initial conditions with no patient variation. For the parameter values in Table 1, model (1) possesses three fixed points: health (green lines), septic death (purple lines), and aseptic death (orange lines). The initial conditions are sampled randomly within the cube:  $0 \leq P_0 \leq 0.42$ ,  $0 \leq N_0 \leq 0.255$ ,  $D_0 = 0$ , and  $0 \leq C_0 \leq 0.35$ . The three fixed points can be differentiated by the steady-state values of  $P$  and  $D$ : health ( $P = 0$ ,  $D = 0$ ), aseptic death ( $P = 0$ ,  $D > 0$ ), and septic death ( $P > 0$ ,  $D > 0$ ). 10, 7, and 3 of the initial conditions evolve to the health, septic death, and aseptic death fixed points, respectively.



**Figure 3.**  $P$ ,  $N$ ,  $D$ , and  $C$  versus time for model (1) with 20 sets of randomly selected parameters with the same initial conditions. 10% patient variation allows the system to reach the health and septic death fixed points with the initial condition  $P_0 = 0.45$ ,  $N_0 = 0.45$ ,  $D_0 = 0$ , and  $C_0 = 0.35$ , whereas only the aseptic death fixed point is obtained for this initial condition with no patient variation. 5, 4, and 11 of the trajectories evolve to the health, septic death, and aseptic death fixed points, respectively.

During the approach to the aseptic death fixed point, the neutrophil level is strong enough to reduce the pathogen level to zero, but the cortisol level is insufficient to reduce the neutrophil level to zero, which leads to increasing damage. Thus, the aseptic death fixed point is characterized by  $P = 0$ ,  $N > 0$ ,  $D > 0$ , and  $C > C_\infty$ .

The outcomes of the immune response to infection can vary from patient to patient, even with the same initial conditions (*e.g.* the pathogen load). To introduce patient variability into the ODE models, we select the parameters ( $\{q\}$ ) in models (1)-(5) randomly from independent Gaussian distributions with mean values  $\mu_q$  in Table 1 and variance  $v$  relative to the mean. Negative values of the parameters can cause the ODE models to become non-integrable, and thus the parameter distributions are cut off so that the parameter values are non-negative. We solve the ODE models for  $10^4$  sets of parameters for each of the  $10^4$  random initial conditions at each  $v$ . The limits for the sampling of the initial conditions for each model are given in Table 3. We then perform a classification analysis on these trajectories to predict the steady-state outcomes. The prediction accuracy  $A$  is defined as the number of correct classifications

of the steady-state outcomes divided by the total number of classifications.

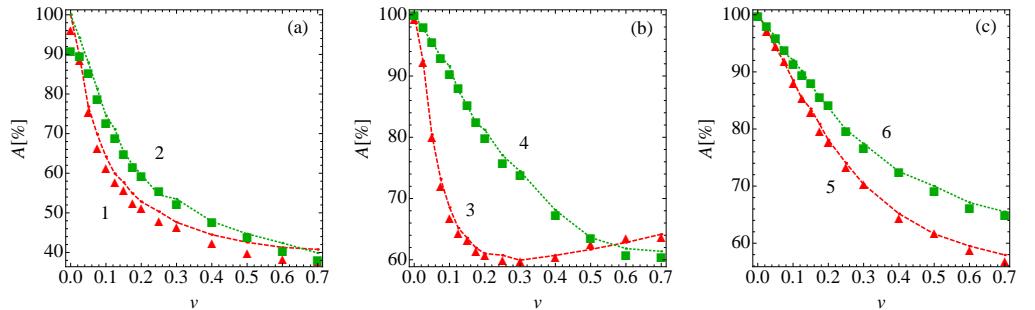
**Table 3. Summary of ODE initial condition ranges**

Model	initial condition ranges
1	$0 \leq P_0 \leq 0.9, 0 \leq N_0 \leq 0.33, D_0 = 0, 0 \leq C_0 \leq 0.5$
2	$0 \leq P_0 \leq 0.1, 0 \leq N_0 \leq 0.15, 0 \leq D_0 \leq 0.1$
3	$0 \leq P_0 \leq 0.2, 0 \leq N_0 \leq 0.3$
4	$0 \leq N_0 \leq 0.15, 0 \leq D_0 \leq 0.1$
5	$0 \leq P_0 \leq 0.7$
6	$0 \leq x_0 \leq 2\pi$

## Results

For a deterministic system of ODEs, the basin of attraction for a given fixed point is defined as the collection of initial conditions that evolve to that particular fixed point. For a given set of parameters, each of the ODE models (1)-(5) possesses well-defined (non-overlapping) basins of attraction for each fixed point.

However, different outcomes can be achieved even for a single initial condition if the parameters of the ODE model are varied. (See Fig. 3.) For example, the ratio of the parameters  $s_c$  and  $\mu_c$  determines the background level of cortisol in model (1). Background cortisol levels are known to vary from patient to patient and can vary from one organ system to another in a given patient. To mimic these variations, we select sets of parameters randomly with mean values in Table 1 and variances  $v$  relative to their mean values. (See Materials and Methods.) With patient variation, an initial condition can possess multiple outcomes, and thus the basins of attraction for the fixed points overlap as shown in Fig. 1.

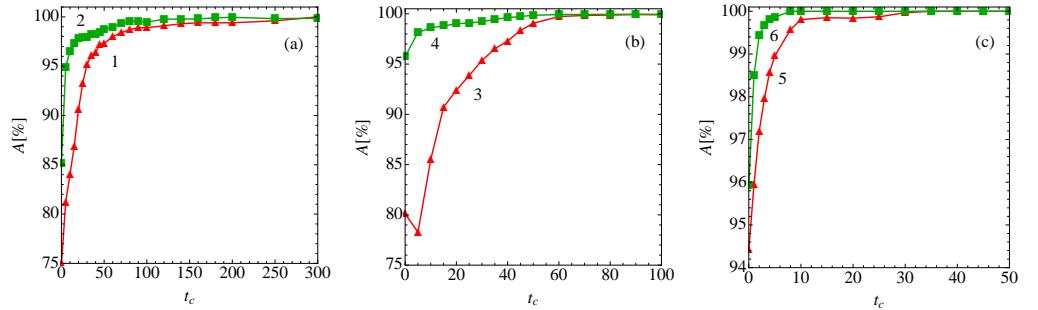


**Figure 4. Prediction accuracy of the steady-state outcome as a function of patient variation.** The prediction accuracy  $A$  using a logistic regression classifier at time  $t_c = 0$  (symbols) and the average best guess over  $10^3$  initial conditions (dashed curves) versus patient variation  $v$  for (a) models (1) and (2), (b) models (3) and (4), and (c) models (5) and (6).

We seek to predict the patient steady-state outcomes in models (1)-(5) in the presence of patient variability  $v$ . For the prediction method, we employ logistic regression with one-versus-all classification [24]. We compare the prediction accuracy at patient variability  $v$  to the average best guess of the steady-state outcome. For the

best guess method, we determine the steady-state outcome for each of  $10^2$  sets of parameters for a given initial condition. We define the best guess as the steady-state outcome with the highest number of occurrences and record the frequency  $f_i$  of the best guess for initial condition  $i$ . We then average the frequency  $f_i$  over  $10^3$  initial conditions for each  $v$  to obtain an estimate for the prediction accuracy in systems with basin overlap.

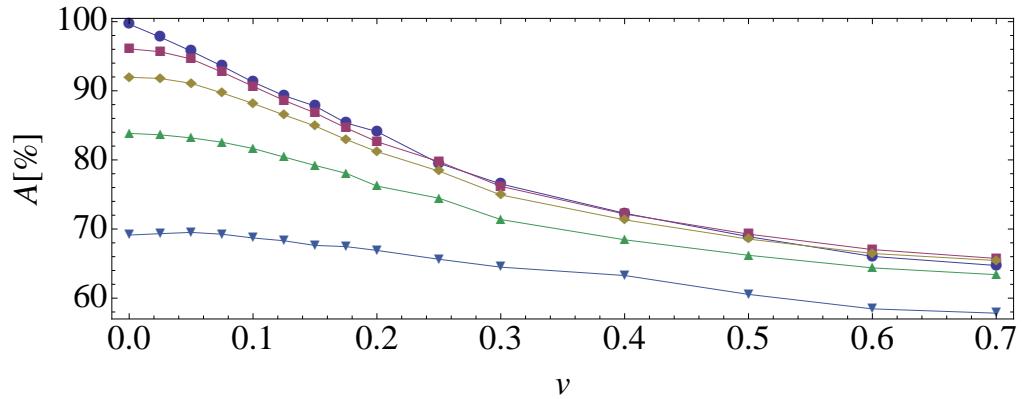
For the prediction method, we solve a given system of ODEs for  $N_i = 10^4$  random initial conditions, each with randomly selected parameter sets with variance  $v$ . We choose  $N_t = 800$  of the  $N_i$  trajectories randomly to train the classifier and predict the outcome of the remaining 9200 trajectories. The classifier maps the state of the system at a given time  $t_c$  to a particular steady-state outcome. The prediction accuracy is then averaged over 10 training and prediction runs, each with  $N_t = 800$  randomly selected training trajectories.



**Figure 5. Prediction accuracy of the steady-state outcome as a function of the classification time.** The prediction accuracy  $A$  using a logistic regression classifier at time  $t_c$  (symbols) for (a) models (1) and (2), (b) models (3) and (4), and (c) models (5) and (6) for patient variation  $v = 0.05$ .

In Fig. 4, we compare the accuracy  $A$  of the logistic regression prediction method (with classification at time  $t_c = 0$ ) to the average best-guess frequency as a function of the patient variability  $v$  for models (1)-(5). For all model ODEs, the prediction accuracy for the logistic regression prediction method is near 100% at  $v = 0$ , decreases for increasing patient variability, and reaches a plateau near  $1/n_f$  in the large  $v$  limit, where  $n_f$  is the number of stable fixed points in the model (except for model (3)). For model (3) with two steady-state outcomes, the prediction accuracy is non-monotonic and increases for  $v > 0.3$  because this ODE model begins to sample parameter regimes where one steady-state outcome is much more probable than the other. In addition, for all models the average best-guess frequency provides an upper bound for the accuracy of the prediction algorithm. Hence, the overlap of the basins of attraction imposes a limit on the prediction accuracy. To test the generality of these results, we studied another one-dimensional ODE (model (6)) with varied fixed point structure compared to that for model (5). (See Supporting Information.) The results for model (6) are very similar to those for models (1)-(5).

In Fig. 4, we showed results for the logistic regression prediction method with classification at  $t_c = 0$ . In Fig. 5, we show the prediction accuracy for models (1)-(6) with patient variability  $v = 0.05$  as a function of the classification time  $t_c$ . For all models, the prediction accuracy grows with increasing  $t_c$ , reaching nearly 100% beyond a characteristic time  $t^*$  that depends on the model. The prediction accuracy improves at later classification times because the system trajectories have moved closer to the fixed points and hence the basins of attraction are more easily separated as shown in Fig. 1 for model (1).



**Figure 6. Prediction accuracy as a function of patient variation for different noise strengths.** The prediction accuracy  $A$  using a logistic regression classifier at time  $t_c = 0$  for model (6) in the presence of measurement noise with strength  $s = 0$  (circles), 0.05 (squares), 0.10 (diamonds), 0.20 (triangles), and 0.50 (triangles).

We also investigated the variation of the prediction accuracy in the presence of measurement noise. We took the trajectories generated for Fig. 4 and added Gaussian random noise to the model variables with variance  $s$  at each time point. We then performed training and testing on the noisy data with classification at time  $t_c = 0$ . In Fig. 6, we show for model (6) that the prediction accuracy decreases with increasing  $s$ . We find similar results for models (1)-(5). These results emphasize that even if the measurement noise could be reduced to zero, the patient variation imposes an intrinsic limitation to outcome prediction.

## Discussion

In clinical settings it is of great importance to determine patient outcomes as quickly as possible with maximum accuracy. In this manuscript, we studied the effects of patient variability on the ability to predict steady-state outcomes in systems of ODEs that model the immune response to infection. For deterministic systems of ODEs with a given fixed set of parameters, each initial condition can be mapped to a given steady-state outcome (or fixed point) and the collection of initial conditions that map to a given steady-state outcome is defined as the basin of attraction of that outcome. Each virtual patient can be defined by a given set of parameters in the model ODE and patient variability can be introduced by varying the model parameters.

We showed that the introduction of patient variation leads to overlaps of the basins of attraction for the steady-state outcomes. In particular, a given initial condition can map to multiple steady-state outcomes for different virtual patients (*i.e.*  $v > 0$ ), which is similar to the case of patients showing different responses to infection in clinical settings. We find that the prediction accuracy of the outcomes decreases strongly with increasing patient variability. Our results emphasize that even when the complete state of the system is known (*i.e.* all patient variables are measured precisely as a function of time), we have limited knowledge of the patient outcome when there is patient-to-patient variability that gives rise to basin overlap.

Our results also show that for all of the model ODEs studied the prediction accuracy increases as the time  $t_c$  used for classification increases. As  $t_c$  increases, the systems move closer to their steady-state outcomes and the basins of attraction separate, which increases the prediction accuracy. Again, this result is consistent with

clinical experience. If a clinician waits to see if the condition of the patient improves or worsens, the prognosis will become more accurate. In our work, we explicitly show that patient-to-patient fluctuations cause a competition between early and accurate outcome prediction.

In this work, we focused on discrete steady-state outcomes (*i.e.* health or death of the patient) of the immune response to infection. However, in many biomedical scenarios, the outcomes involve continuous variables rather than discrete states. In future work, we will apply similar techniques to understand the effects of patient variability on the predictions of continuous model variables, for example, the immune response and vaccination efficacy for influenza [25].

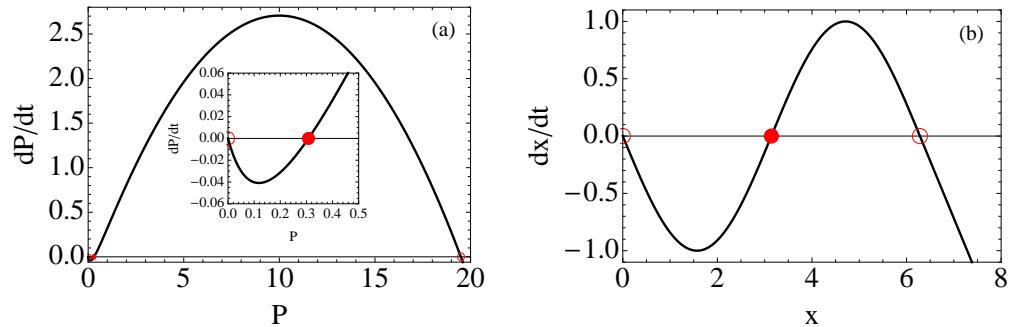
## Supporting Information

In this section, we describe a generalized one-dimensional ODE that has the same fixed point structure as model (5) (Eq. 15), but different locations for the fixed points. We also provide technical details for the logistic regression one-versus-all classifier employed in this work.

Model (6) is a one-dimensional ODE for the variable  $x$ :

$$dx/dt = \begin{cases} B \cos(kx + \pi/2), & \text{if } 0 \leq x \leq 2\pi/k \\ -B k \sin(5\pi/2)(x - 2\pi/k), & \text{if } 2\pi/k < x \end{cases} \quad (16)$$

with two parameters  $k = 1$  and  $B = 1$ . Model (5), which is a one-dimensional ODE for pathogen  $P$ , possesses three fixed points:  $P = 0$ , 0.3078, and 19.49 for the mean parameters in Table 1. As shown in Fig. 7 (a), the two outer fixed points for model (5) are stable, and the middle fixed point, which is near zero, is unstable. For model (6), the central unstable fixed point is moved to the midpoint of the two outer stable fixed points and the shape of the function is changed to retain the fixed point structure.



**Figure 7. Comparison of ODE models (5) and (6).** The functions  $f(P) = dP/dt$  and  $f(x) = dx/dt$  for models (a) (5) and (b) (6). (See Eqs. 15 and 16.) Stable and unstable fixed points are marked by open and filled circles, respectively. Both models share the same fixed point topology: one stable fixed point at zero and one at a positive value. The unstable fixed point lies between the two stable fixed points. The inset in (a) magnifies  $dP/dt$  near the origin.

For the prediction method, we train a logistic regression classifier in a one-versus-all scheme to predict the steady-state outcomes based on the values of the set of variables  $x^{(i)} = \{x_1^i, x_2^i, \dots, x_N^i\}$  at a specified time  $t_c$ , where  $N$  is the number of variables in the model. To distinguish between two outcomes with logistic regression,

we label the outcomes,  $y^{(i)} = 0$  and  $y^{(i)} = 1$  for the  $i$ th set of variables. We fit a logistic (or sigmoidal) function to the input data such that the function

$$P_\theta(x^{(i)}) = \frac{1}{1 + e^{\theta_0 + \theta_1 x_1^i + \theta_2 x_2^i + \dots + \theta_N x_N^i}} \quad (17)$$

gives the probability that  $y^{(i)} = 1$  given the input data  $x^{(i)}$ . The parameters  $\theta_j$ , where  $j = 1, 2, \dots, N$ , are determined by minimizing the cost function

$$J(\theta_1, \theta_2, \dots, \theta_N) = -\frac{1}{m} \left[ \sum_{i=1}^m y^{(i)} \log(P_\theta(x^{(i)})) + (1 - y^{(i)}) \log(1 - P_\theta(x^{(i)})) \right], \quad (18)$$

where  $m$  is the number of training samples. Since models (1) and (2) possess three steady-state outcomes (aseptic death, septic death, and health), we must go beyond the binary classification scheme described above. To classify ODE models with three steady-state outcomes, we implement the one-versus-all classification scheme. To do this, we consider three labeled outcomes,  $y_h^{(i)}$ ,  $y_{ad}^{(i)}$ , and  $y_{sd}^{(i)}$ , for a given set of variables  $x^{(i)}$ .  $y_h^{(i)} = 0$  if the patient outcome is not health (*i.e.* aseptic or septic death) and  $y_h^{(i)} = 1$  if the patient outcome is health. Similar definitions apply for  $y_{ad}^{(i)}$  and  $y_{sd}^{(i)}$ . (See Table 4.) We use these labeled outcomes and Eq. 18 to determine  $P_h(x)$ ,  $P_{ad}(x)$ , and  $P_{sd}(x)$ . Given an unlabeled set of variables ( $x = \{x_1, x_2, \dots, x_N\}$ ), we calculate  $P_h$ ,  $P_{ad}$ , and  $P_{sd}$  and select the outcome with highest probability to be the predicted outcome.

**Table 4. One-versus-all outcomes**

Class 1 ( $y = 0$ )	Class 2 ( $y = 1$ )
septic death + health	aseptic death
health + aseptic death	septic death
aseptic death + septic death	health

## Acknowledgments

This work was partially supported by DARPA (Space and Naval Warfare System Center Pacific) under award number N66001-11-1-4184 and partially supported by the National Science Foundation Grant No. NSF PHY11-25915. Additional support for this work was provided by the Infectious Disease Supercluster, Colorado State University, 2012 seed grant. This work also benefited from the facilities and staff of the Yale University Faculty of Arts and Sciences High Performance Computing Center and NSF Grant No. CNS-0821132 that partially funded acquisition of the computational facilities.

## References

1. Koup, R. A., Safrit, J. T., Cao, Y., Andrews, C. A., McLeod, G., Borkowsky, W., Farthing, C., & Ho, D. D. (1994). Temporal association of cellular immune responses with the initial control of viremia in primary Human Immunodeficiency Virus Type 1 Syndrome. *J Virol.* 68(7), 4650-4655.
2. Guidotti, L. G., & Chisari, F. V. (2001). Noncytolytic control of viral infections by the innate and adaptive immune response. *Annu Rev Immunol.* 19(1), 65-91.
3. Janeway Jr, C. A., & Medzhitov, R. (2002). Innate immune recognition. *Annu Rev Immunol.* 20(1), 197-216.
4. Klipp, E. (2005). *Systems Biology in Practice: Concepts, Implementation And Application.* Wiley-VCH, Weinheim.
5. Masopust, D., Vezys, V., Marzo, A. L., & Lefrancois, L. (2001). Preferential localization of effector memory cells in nonlymphoid tissue. *Science*, 291(5512), 2413-2417.
6. Pope, C., Kim, S. K., Marzo, A., Williams, K., Jiang, J., Shen, H., & Lefrancois, L. (2001). Organ-specific regulation of the CD8 T cell response to *Listeria monocytogenes* infection. *J Immunol.* 166(5), 3402-3409.
7. Young, D., Stark, J., & Kirschner, D. (2008). Systems biology of persistent infection: Tuberculosis as a case study. *Nat Rev Microbiol.* 6(7), 520-528.
8. Perelson, A. S., & Nelson, P. W. (1999). Mathematical analysis of HIV-1 dynamics in vivo. *SIAM Rev.* 41(1), 3-44.
9. Callaway, D. S., & Perelson, A. S. (2002). HIV-1 infection and low steady state viral loads. *B Math Biol.* 64(1), 29-64.
10. Nelson, P. W., & Perelson, A. S. (2002). Mathematical analysis of delay differential equation models of HIV-1 infection. *Math Biosci.* 179(1), 73-94.
11. Perelson, A. S. (2002). Modeling viral and immune system dynamics. *Nat Rev Immunol.* 2(1), 28-36.
12. Bocharov, G. A., & Romanyukha, A. A. (1994). Mathematical model of antiviral immune response III. Influenza A virus infection. *J Theor Biol.* 167(4), 323-360.
13. Belz, G. T., Wodarz, D., Diaz, G., Nowak, M. A., & Doherty, P. C. (2002). Compromised influenza virus-specific CD8+-T-cell memory in CD4+-T-cell-deficient mice. *J Virol.* 76(23), 12388-12393.
14. Dowdy, D. W., Chaisson, R. E., Moulton, L. H., & Dorman, S. E. (2006). The potential impact of enhanced diagnostic techniques for tuberculosis driven by HIV: A mathematical model. *Aids* 20(5), 751-762.
15. Wodarz, D., & Nowak, M. A. (2002). Mathematical models of HIV pathogenesis and treatment. *BioEssays* 24(12), 1178-1187.
16. Serdoba, I., & Kieny, M. P. (2006). Assembling a global vaccine development pipeline for infectious diseases in the developing world. *Am J Public Health.* 2006 September; 96(9): 1554-1559.

17. Wetterstrand KA. DNA Sequencing Costs: Data from the NHGRI Genome Sequencing Program (GSP) Available at: [www.genome.gov/sequencingcosts](http://www.genome.gov/sequencingcosts). Accessed Jan. 2015.
18. Day, J., Rubin, J., Vodovotz, Y., Chow, C. C., Reynolds, A., & Clermont, G. (2006). A reduced mathematical model of the acute inflammatory response II. Capturing scenarios of repeated endotoxin administration. *J Theor Biol.* 242(1), 237-256.
19. Vodovotz, Y., Csete, M., Bartels, J., Chang, S., & An, G. (2008). Translational systems biology of inflammation. *PLoS Comput Biol.* 4(4), e1000014.
20. Vodovotz, Y., Constantine, G., Rubin, J., Csete, M., Voit, E. O., & An, G. (2009). Mechanistic simulations of inflammation: Current state and future prospects. *Math Biosci.* 217(1), 1-10.
21. Wang, K., Bhandari, V., Chepustanova, S., Huber, G., Stephen, O., Corey, S. O., M.D. Shattuck, & Kirby, M. (2013). Which biomarkers reveal neonatal sepsis?. *PLoS One*, 8(12), e82700.
22. O'Hara, S., Wang, K., Slayden, R. A., Schenkel, A. R., Huber, G., O'Hern, C. S., M.D. Shattuck, & Kirby, M. (2013). Iterative feature removal yields highly discriminative pathways. *BMC Genomics* 14(1), 832.
23. Reynolds, A., Rubin, J., Clermont, G., Day, J., Vodovotz, Y., & Bard Ermentrout, G. (2006). A reduced mathematical model of the acute inflammatory response: I. Derivation of model and analysis of anti-inflammation. *J Theor Biol.* 242(1), 220-236.
24. Hastie, T., Tibshirani, R., Friedman, J., Hastie, T., Friedman, J., & Tibshirani, R. (2009). *The elements of statistical learning* (Vol. 2, No. 1). New York: Springer.
25. Hancioglu, B., Swigon, D., & Clermont, G. (2006). A dynamical model of human immune response to Influenza A virus infection. *J Theor Biol.* 246, 70-86.